

GROUND DELAY PROGRAMS: COLLABORATIVE PLANNING UNDER UNCERTAINTY

Authors:

Avijit Mukherjee

Associate Project Scientist

U.C. Santa Cruz

Building 19, Moffett Field

Mountain View, CA, 94035

Phone: 650-604-0380

Email: avijit@ucsc.edu

Mark Hansen

Professor

Department of Civil and Environmental Engineering

107 McLaughlin Hall, U.C. Berkeley, CA 94720

Phone: 510-642-2880

Email: mhansen@ce.berkeley.edu

ABSTRACT

Ground delay programs (GDPs) are implemented by the Federal Aviation Administration (FAA) to mitigate arrival demand – capacity imbalance at an airport. Typically such conditions occur due to the impact of adverse weather on airport capacity. During a GDP flights are assigned arrival slots which may be later than their scheduled time of arrival. Under the Collaborative Decision Making (CDM) paradigm, which has been adopted since mid 90's, airlines are allowed to perform substitutions among the slots they have been assigned in order to reduce their internal costs and achieve their business objectives. Airlines may also cancel some of their flights and utilize the vacant slots by substituting some of their own flights. However, in certain circumstances, an airline may not be able to substitute one of its flights to such a slot, maybe because the time stamp associated with the slot is too early for substituting any flight. Therefore, after a round of intra-airline substitutions, it is possible that some arrival slots remain unutilized. The FAA then runs the Compression Algorithm to re-allocate arrival slots among flights that are not cancelled by airlines. The basic principle behind this algorithm is that whenever possible a delayed flight is assigned an earlier arrival slot, so that the unused capacity during any time period is reduced (or eliminated if possible). While reallocating a flight to an open slot, preference is given to flights that belong to the airline which has released the slot.

In reality, GDPs are implemented in face of uncertainty in airport capacity. The uncertainty arises from imperfect knowledge on how the weather might evolve in the later time periods. In practice, however, the FAA assumes a deterministic forecast of arrival capacity and implements GDP based on that. In the recent years, there has been some research on developing optimization models to plan GDPs under probabilistic forecasts.

In this paper, we present models that will enable collaborative decision making while planning GDPs in face of uncertainty. We first describe two recently developed stochastic optimization models for planning GDPs. The first is a static model, which means ground delay decisions are made once at the beginning of planning horizon and not updated later. In the second model, decisions are dynamically revised based on updated information on airport capacity. We then present a methodology which will allow airlines to perform intra-airline substitutions using the solutions from the stochastic models for planning GDPs. Finally, we present a model that performs analogous to the CDM Compression algorithm, which can be applied after the round of airline specific substitutions and cancellations. The objective of this model is to efficiently utilize the vacant slots that become available after cancellations.

1. INTRODUCTION

Meteorological conditions, particularly convective weather, can significantly impact the capacity of an airport. Other forms of adverse weather such as fog, snow/ice, low cloud ceiling and poor visibility can also reduce airport capacities significantly. Reduction in airport capacity may induce delays among flights. Typically, under reduced capacity conditions at an airport, the Federal Aviation Administration (FAA) implements a Ground Delay Program (GDP), in which certain flights are assigned pre-departure delays (or ground delays) at their respective origin airports. This is done to mitigate the demand-capacity imbalance at the destination airport caused by weather. If weather conditions become more severe than anticipated, some aircraft may also face airborne delay while is en-route to the destination airport. Airborne delay, although quite common in the National Airspace System (NAS) during adverse weather conditions, is undesired mainly due to expensive fuel consumption and safety issues related to controlling aircraft and maintaining separation standards while they are airborne. This has motivated researchers, since late 80's, to develop optimization models and algorithms to decide on the amount of ground delays that must be imposed on various flights when there is capacity shortfall at one or more airports. The problem of assigning ground delays to various flights subject to the airport capacity constraints in order to minimize an objective function, which is typically the sum of ground and airborne delays, weighted by their relative costs, is known as the ground holding problem (see Odoni (1987) for a systematic description of the ground holding problem).

Within the domain of the ground holding problem, there are two sub-problems: the single airport ground holding problem (SAGHP) and the multi-airport ground holding problem

(MAGHP). In the SAGHP, the problem is solved for one destination airport at a time. In the MAGHP, a network of airports is considered, so that delay on a given flight segment can propagate to downline segments flown by the same aircraft. In this paper, we will focus on the SAGHP.

As theoretical research on ground holding problems has progressed, air traffic flow management has evolved from a completely centralized system to one in which users have more autonomy about how to adapt their schedules when adverse conditions reduce airport capacity. At the first stage of implementing a GDP, the FAA allocates arrival slots to airlines based on their original flight schedules and the first-come/first-served principle, commonly known as the ration-by-schedule (RBS) allocation. To counter uncertainty in airport capacity during future time intervals, which may change depending on how the weather materializes, the FAA might also exempt certain flights, which originate from airports that are located beyond certain geographical distance from the destination, from the GDP. The exempt flights are allowed to arrive at their scheduled times. After the initial round of slot allocations, airlines are allowed to perform substitutions among the slots they have been assigned in order to reduce their internal costs and achieve their business objectives. Airlines may also cancel some of their flights and utilize the vacant slots by substituting some other flights they own and operate. However, in certain circumstances, an airline may not be able to substitute one of its flights to such a slot, maybe because the time stamp associated with the slot is too early for substituting any flight. Therefore, after a round of intra-airline substitutions, it is possible that some arrival slots remain open. The FAA then runs the Compression Algorithm to re-allocate arrival slots among flights that are not cancelled. The basic principle behind this algorithm is that whenever possible a delayed flight is assigned an earlier arrival slot, so that the unused capacity during any time period is reduced (or eliminated if possible). While

reallocating a flight to an open slot, preference is given to flights that belong to the airline which has released the slot. Such collaborative planning of ground delay programs has resulted from adopting the Collaborative Decision Making (CDM) philosophy in the air traffic flow management world since mid 90's (see Hoffman 1997, and Vossen 2002 for detailed discussion on CDM and its implications on planning GDPs in real world).

In this paper, we present models that will enable collaborative decision making while planning GDPs in face of uncertainty in airport capacity. We first describe two recently developed stochastic optimization models for the SAGHP. The first is a static model, which means ground delay decisions are made once at the beginning of planning horizon and not updated later. In the second model, decisions are dynamically revised based on updated information on airport capacity. We then present a methodology which will allow collaborative planning of ground delay programs in face of uncertainty using these stochastic SAGHP models. We describe how the FAA can use these models to do initial slot allocation to airlines, who can then perform intra-airline substitutions and cancellations.

2. BACKGROUND

In a deterministic setting, where the airport capacities during future time intervals are known in advance with perfect information, the SAGHP can be formulated as a minimum-cost network flow problem (Tarrab and Odoni, 1993); although, there are certain variants of the deterministic SAGHP that are computationally harder (Hoffman and Ball 2000). In reality, however, weather forecasts are rarely perfectly accurate. In a stochastic setting, where airport capacities in future time intervals are uncertain, the formulation of the SAGHP becomes complex. There is a moderate amount of literature on stochastic optimization models for the

SAGHP. In most of the models, the objective is to minimize the expected total cost of delay. The optimization models are of two types: static stochastic, and dynamic stochastic. Static (or single stage) stochastic models account for uncertainty, but do not utilize updated information on the evolution of airport capacity, while the dynamic model accounts for the ability to re-assign ground delays depending on what weather conditions unfold.

Richetta and Odoni (1993) first proposed an integer programming model to solve the multi-period static stochastic SAGHP. In their model, uncertainty in airport capacity is represented by a finite set of scenarios, each of which represents a time-varying profile of the airport capacity that is likely to occur. Recently, Kotnyek and Richetta (2006) showed that when the ground holding costs are marginally increasing the dual of the Richetta and Odoni (1993) model becomes a minimum cost network flow problem in formulation, and hence integer solutions are guaranteed to obtain by solving the LP relaxation of the model.

Ball et al. (2003) proposed another variant of the static stochastic SAGHP formulated as an integer-programming (IP) model. The Ball et al. model has the following important properties. The constraint matrix of the formulation is totally unimodular, which means integer solutions are guaranteed from LP relaxation of the model. Therefore the model can be applied in real time to solve practical scale problems. The Ball et al. model do not decide on how much ground delay to impose on individual flights. Rather, the decision variables – the planned airport acceptance rates (PAARs) -- are the aggregate number of aircraft planned to arrive at the airport during various time intervals. In practice, when implementing a ground delay program, the FAA first sets the airport acceptance rates (the vernacular for the maximum number of arrivals that the airport can accommodate) during different time intervals, and then allocates arrival slots to airlines using certain rationing schemes, described

later in this paper. Hence the Ball et al. model can be used in real world to set the optimum set of airport acceptance rates.

Quite frequently, in practice, traffic managers make adjustments to a ground delay program to compensate for unanticipated change in airport capacity. This served as a motivation behind the development of dynamic stochastic models for the SAGHP, which are capable of revising ground delay decisions based on evolving forecasts. Richetta and Odoni (1994) formulated a (partially) dynamic multi-stage stochastic IP model. Rather than assigning delays to all flights at once, the Richetta-Odoni model assigns delay as the scheduled departure time approaches, so that decisions can be made with the most up-to-date forecast information. Uncertainty in airport arrival capacities is represented through a finite number of scenarios arranged in a probabilistic binary decision tree. As the day progresses, branches of the tree are realized, resulting in better information about future capacities. In Richetta-Odoni model, however, ground delays, once assigned, cannot be revised, even though this is technically possible so long as the flight has not yet departed.

More recently, Mukherjee and Hansen (2003) developed a dynamic multi-stage stochastic IP formulation for the SAGHP, in which ground delays assigned to various flights can be revised during different decision stages, based on updated information on airport operating conditions. An additional advantage of the Mukherjee-Hansen model is that it can handle a wide range of linear and non-linear cost functions in the objective. For example, in addition to the standard linear delay cost function with different weights for airborne and ground delay, we can minimize expected squared arrival delay against schedule, and expected squared deviation between the assigned arrival slot and the ideal RBS arrival slot. Multiple objective functions can also be considered through weighting schemes that either balance

different objectives or introduce subsidiary objectives as “tie-breakers” when the original model has multiple optima.

In all the stochastic SAGHP models discussed above, uncertainty in airport conditions is captured through a finite set of scenarios, each representing a possible evolution of the time-varying profile of the airport capacity. The dynamic models require more than just the set of capacity scenarios and their unconditional probabilities. In both Richett-Odoni (1994) and Mukherjee-Hansen (2003) models, a probabilistic scenario tree is required, whose branching reveals updated information on the evolving conditions. Recently, Liu et al. (2006) developed a methodology for constructing capacity scenarios, using statistical clustering techniques, and scenario trees from the recorded data on airport acceptance rates. Figure 1 below shows six possible capacity profiles, each of which is likely to occur at San Francisco Intl. airport (SFO) on any given day. Figure 2 shows the corresponding scenario tree, whose branching points reveal which scenario or a set of scenarios are likely to prevail.

3. STOCHASTIC MODELS FOR THE SAGHP

In this section we review the two recently developed stochastic models for the SAGHP: (1) the Ball et al. (2003) static model, and (2) the Mukherjee and Hansen (2003) dynamic model. We first present the formulations and then compare the solutions obtained by applying these models for a realistic problem.

3.1. The Ball et al. (2003) Static Stochastic Model for the SAGHP

As in most of the discrete optimization models, the planning horizon is divided into T equal time-periods. Let there be Q capacity scenarios, each scenario depicting a possible evolution of airport arrival capacity over the planning period with the scenario q ($q = 1, 2, \dots, Q$) having a probability of occurrence equal to p_q . Let M_t^q denote the capacity at time period t under the scenario q . In order to ensure that all flights that are scheduled to land do get assigned to land during a time period, let there be a time period $T+1$ with unlimited capacity under all scenarios. Let N_t be the number of flights scheduled to arrive during time interval t ($t = 1, 2, \dots, T+1$). The PAARs are the decision variables, and are denoted by A_t ($t = 1, \dots, T+1$). Let G_t denote the number of flights that are delayed on ground from time period t to $t+1$. If the numbers of aircraft that arrive at the terminal airspace exceed airport capacity, certain flight might face airborne holding. Let W_t^q denote the number of aircraft that are unable to land during time period t under scenario q , and hence faces airborne holding during that time period. Let λ denote the cost ratio between unit airborne and ground delay of any flight.

The objective function and the set of constraints are defined as follows:

$$\text{Min. } \sum_{t=1}^T G_t + \sum_{q=1}^Q p_q \sum_{t=1}^T \lambda W_t^q \quad (1)$$

s.t.

$$A_t - G_{t-1} + G_t = N_t; t = 1, \dots, T+1; (G_0 = G_{T+1} = 0) \quad (2)$$

$$A_t + W_{t-1}^q - W_t^q \leq M_t^q; t = 1, \dots, T+1; q = 1, \dots, Q, (W_0^q = W_{T+1}^q = 0) \quad (3)$$

$$A_t, G_t, W_t^q \geq 0 \text{ and integer} \quad (4)$$

The objective function minimizes the weighted sum of expected ground and airborne delays. Constraint set (2) is required for flow conservation, whereas constraints (3) ensures that the number of aircraft that actually land during any time interval under a scenario do not exceed

the corresponding airport capacity. The solution to the model yields the optimum values of the decision variables -- A_t -- that reflect the number of aircraft that are planned to arrive, not necessarily land, at the terminal airspace and desire to land during various time intervals.

3.2. Mukherjee and Hansen (2003) Dynamic Stochastic Model for the SAGHP

Wherever possible, we use the same notations as in the static stochastic model described above. As mentioned before, in addition to the capacity scenarios and their unconditional probabilities, the dynamic model requires as input how the scenario tree unfolds. The following parameters are used to provide the scenario tree information (see Mukherjee, 2004, and Liu et al., 2006 for details). Let B be the total number of branches in the scenario tree; $B \geq Q$. Each branch corresponds to a set of scenarios. The π_b scenarios corresponding to branch $b \in \{1, \dots, B\}$ is given by the set $\Omega_b = \{S_1^b, \dots, S_k^b, \dots, S_{\pi_b}^b\}; S_k^b \in \{1, \dots, Q\}$. The time periods corresponding to start and end nodes of branch b are denoted by o_b and $\mu_b; b \in \{1, \dots, B\}$. The other input parameters are the set of flights $\Phi = \{1, \dots, F\}$, their scheduled departure and arrival time periods -- a_f and d_f respectively, scenario-specific airport capacities M_t^q , $q = \{1, \dots, Q\}$, and the cost ratio between unit airborne and ground delay -- λ .

The decision variables X_{ft}^q ($q \in \{1, \dots, Q\}, f \in \Phi, t \in \{a_f, \dots, T+1\}$) are binary, and are defined as follows:

$$X_{ft}^q = \begin{cases} 1 & \text{if flight } f \text{ is planned to arrive by the end of time period } t \\ & \text{under scenario } q; \\ 0 & \text{otherwise} \end{cases}$$

Corresponding to the variables $X_{f,t}^q$ is a set of corresponding auxiliary variables for the departure time period, defined as follows:

$$Y_{f,t}^q = \begin{cases} 1 & \text{if flight } f \text{ is released for departure by the end of time period } t \\ & \text{under scenario } q; \\ 0 & \text{otherwise} \end{cases}$$

The departure release variables track the planned arrival times but are displaced earlier in time by the amount $a_f - d_f$. Hence the variables $Y_{f,t}^q$ are related to $X_{f,t}^q$ as follows:

$$Y_{f,t}^q = \begin{cases} X_{f,t+a_f-d_f}^q; & \text{if } t+a_f-d_f \leq T \\ 1 & \text{otherwise} \end{cases} \quad (5)$$

The objective function and the set of constraints are given as follows:

$$\text{Min } \sum_{q=1}^Q p_q \left(\sum_{f \in \Phi} \sum_{t=Arr_f}^{T+1} (t - Arr_f) (X_{f,t}^q - X_{f,t-1}^q) + \lambda \sum_{t=1}^T W_t^q \right) \quad (6)$$

s.t:

$$X_{f,t}^q - X_{f,t-1}^q \geq 0; \quad f \in \Phi, q \in \{1, \dots, Q\}, t \in \{a_f, \dots, T+1\}; (X_{f,T+1}^q = 1) \quad (7)$$

$$W_{t-1}^q - W_t^q + \sum_{f \in \Phi} (X_{f,t}^q - X_{f,t-1}^q) \leq M_t^q; \quad t \in \{1, \dots, T+1\}, q \in \{1, \dots, Q\}; (W_0^q = W_{T+1}^q = 0) \quad (8)$$

$$Y_{f,t}^{S_1^b} = \dots = Y_{f,t}^{S_k^b} = \dots = Y_{f,t}^{S_{\pi_b}^b}; \quad f \in \Phi, t \in \{1, \dots, T\}; S_k^b \in \Omega_b: \pi_b \geq 2 \text{ and } o_b \leq t \leq \mu_b \quad (9)$$

As in most of the models for SAGHP, the objective function minimizes the expected total delay cost. Constraint set (7) ensures that the decision variables $X_{f,t}^q$ are non-decreasing. Constraints (8) are similar to constraints (3) described above. Constraints (9) are the set of “coupling” constraints, sometimes known as non-anticipatory constraints in the literature (see Birge and Louveaux, 1997), on the ground-holding decision variables $Y_{f,t}^q$. These constraints

force ground delay decisions to be made solely on information available at time t . For a given time period t , it is required that the ground holding decisions are the same for all scenarios associated with the same scenario tree branch b (in other words the scenarios belonging to the set Ω_b) in that time period.

3.3. Static vs. Dynamic SAGHP

We apply the above models – the Mukherjee and Hansen (2003) and the Ball et al. (2003) – to a realistic problem, and compare the solutions. The FAA – ASPM¹ database was used to obtain flight schedules at the Dallas Fort Worth Intl. airport (DFW) on a summer weekday of 2003. A hypothetical ground holding problem was constructed, described as follows. It is assumed that weather impacts the airport capacity by reducing it to 60 arrival operations per hour during the morning hours. There are six capacity scenarios, each corresponding to a possible weather clearance time between 8:00AM and 10:30AM in 30 minutes increment. When weather clears, the capacity of the airport rises to its visual-flight-rule value, which is 135 arrivals/hour. Figure 3 shows the capacity profiles for the six capacity scenarios $q = 1, \dots, 6$. The scenario probabilities are assumed as follows: $p_q = (0.4, 0.2, 0.1, 0.1, 0.1, 0.1)$.

Table 1 shows the expected ground and airborne delays and the total delay costs that are obtained when the Ball et al. (2003) static and the Mukherjee and Hansen (2003) dynamic models are applied, under two different cost ratios $\lambda = 3$ and $\lambda = 25$, to the above problem. When the cost ratio is relatively low – i.e. $\lambda = 3$ -- airborne holding is faced in both models. The static model assigns relatively low ground delay compared to the dynamic model and releases some flights placing a bet on the capacity to improve. The dynamic model on the

¹ www.apo.data.faa.gov/

other hand adopts a “wait-and-see” policy and assigns relatively more ground holding initially. If capacity improves, ground delays are revised and flights are released earlier, otherwise, ground delays are extended. The adaptive policy of assigning ground delays leads to less risk of facing airborne holding in the dynamic model. When the cost ratio is high – i.e. $\lambda = 25$ -- both models adopt a conservative ground holding policy. The dynamic model achieves lower expected ground delays by virtue of releasing flights if conditions improve. Such dynamic revision capability is absent in the static model, which faces about 42% more total delay cost compared to the dynamic model.

4. COLLABORATIVE GDP PLANNING UNDER UNCERTAINTY

The stochastic models described above can be used as decision support tools to plan ground delay programs in face of uncertainty. However, in order to implement them in practice, it is necessary to provide a methodology which will allow the airlines to make intra-airline substitutions after the FAA assigns slots to them using the solutions from either model. In this section, we discuss how the Ball et al. (2003) static and the Mukherjee-Hansen (2003) dynamic stochastic models for the SAGHP can be used to plan GDPs while allowing collaborative decision making involving the users (the airlines) and the controlling entity (the FAA).

The collaborative GDP planning using the Ball et al. (2003) static model is described through Figure 4 below. As mentioned before, the Ball et al. model produces the optimal set of “planned airport acceptance rates” (PAARs). The FAA can use these numbers to assign slots to the airlines using any rationing scheme of their choice. For example, geographical exemptions followed by ration-by-schedule algorithm can be used for the initial slot

allocation. Airlines can then perform substitutions and cancellations, and provide their updated schedules/flight-plans to the FAA, who can then run the Compression algorithm to re-allocate slots among airlines while utilizing the vacant slots created by some of the cancelled flights. An important property of the Ball et al. model is that its objective function remains invariant under intra-airline substitution. As time progresses, updated information on capacity and demand becomes available, and the whole process can be re-executed in order to be adaptive to the evolving conditions.

The solution to the Mukherjee-Hansen (2003) model, although assigning slots to flights, may in practice be used as a means of allocating slots to airlines. The arrival slots are however, contingent upon which scenario (or a bundle of scenarios) is realized at the time a flight is released. The planned arrival time of a flight $f \in \Phi$ under a scenario q is given by the

expression $\sum_{t=a_f}^{T+1} t(X_{f,t}^q - X_{f,t-1}^q)$. We denote this value by θ_f^q . Therefore for each flight, the

model solves for an optimal portfolio of scenario contingent arrival times. For any given flight and a portfolio of scenario contingent arrival time periods (or slots) it is possible to perform a feasibility check of assigning the portfolio of slots to the flight. Feasibility depends on the flight's scheduled departure time and the coupling constraints given by expression (9).

Definition 1: The assignment of flight f to the portfolio of scenario-contingent arrival time periods $\tau^q, q \in \{1, \dots, Q\}$ is feasible, iff $\tau^q \geq a_f$, and the variables $Y_{f,t}^q \geq d_f$ satisfy the coupling constraints (9) defined in Section 3.2.

Any substitution involving two flights requires swapping the scenario specific arrival slots between them. Therefore, for the substitution to occur, each flight must be assigned a new portfolio of slots that was owned by the other flight.

Definition 2: A substitution of scenario contingent slots between any two flights f_1 and f_2 is allowed if $\theta_{f_1}^q$ ($q \in \{1, \dots, Q\}$) is a feasible assignment for f_2 , and $\theta_{f_2}^q$ is feasible for f_1 .

Scenario based substitution interchanges arrival time periods between two flights under different scenarios, and therefore doesn't affect the arrivals (and hence ground delays) of other flights that are not involved in the substitutions. Also, the scenario-specific ground delays are not affected by a substitution of scenario-contingent arrival times between two flights. This is proved in the following proposition.

Proposition 1: A scenario contingent slot substitution between two flights does not change the total ground delays under different scenarios.

Proof: Let f_1 and f_2 be two flights involved in substitution. Let scenario-specific ground delays of any flight $f \in \Phi$ before substitution be denoted by $g_f^q, q \in \{1, \dots, Q\}$. After substitution, the scenario-specific ground delays of flight $f \in \Phi \setminus \{f_1, f_2\}$ remains unchanged. For flights f_1 and f_2 the ground delays after substitution are given by $g_{f_1}^q + \theta_{f_2}^q - \theta_{f_1}^q$ and $g_{f_2}^q + \theta_{f_1}^q - \theta_{f_2}^q$ respectively. Therefore the scenario-specific total ground delays of the two flights involved in the substitution remains unchanged. The proof follows.

Proposition 2: Scenario-specific planned arrival rates during any time period $t \in \{1, \dots, T+1\}$ remain unchanged after substitutions.

Proof: Let f_1 and f_2 be two flights involved in a substitution. Scenario-specific arrival time periods of flights $f \in \Phi \setminus \{f_1, f_2\}$ – i.e. $X_{f,t}^q$ – remain unchanged after substitution.

Therefore for time periods $t \in \{1, \dots, T+1\} \setminus \{\theta_{f_1}^q \cup \theta_{f_2}^q\}$ scenario-specific planned arrival numbers remains unchanged. Before substitution the expression $(X_{f_1,t}^q - X_{f_1,t-1}^q)$ attains a value 1 for $t = \theta_{f_1}^q$ and 0 otherwise. Similarly before substitution $(X_{f_2,t}^q - X_{f_2,t-1}^q) = 1$ iff $t = \theta_{f_2}^q$. After substitution, $(X_{f_1,t}^q - X_{f_1,t-1}^q) = 1$ iff $t = \theta_{f_2}^q$ and $(X_{f_2,t}^q - X_{f_2,t-1}^q) = 1$ iff $t = \theta_{f_1}^q$. Therefore during time periods $t \in \{\theta_{f_1}^q \cup \theta_{f_2}^q\}$, the scenario specific planned arrivals after substitution remains same as that before. Q.E.D.

Corollary 1: Scenario-specific airborne delays remain same after substitution.

Proof: From proposition 2, during any time period, the number of planned arrivals --

$\sum_{f \in \Phi} (X_{f,t}^q - X_{f,t-1}^q)$ -- remains the same before and after substitution. Therefore, it is easy to

see, from expression (8), that the time-varying airborne queueing delay – W_t^q – remains unchanged.

Corollary 2: The value of the objective function of the Mukherjee-Hansen model, given in expression (6), remains invariant to a feasible substitution.

Proof. The proof follows from Proposition 1 and Corollary 1.

The above discussion focused on scenario-contingent slot substitution between two flights at a time. However, an airline may want to perform substitutions involving multiple flights. For example, a flight f_1 may be substituted for f_2 under a particular scenario, and for f_3 under a different scenario. Furthermore, an airline may cancel some of its flights and utilize the vacant slots to reduce delays of some of its other flights. From the solutions to the Mukherjee-Hansen model it is possible to determine the scenario-specific numbers of slots -- v_{at}^q -- assigned to each an airline $a \in A$ during a time interval t , where A is the set of airlines. The FAA allocates these scenario-contingent slots to airlines. The airlines can then assign a ground delay cost function to individual flights they own, and re-allocate slots among flights. This is similar to present day substitutions performed in a ground delay program. However, in order to preserve feasibility, airlines must abide by the coupling constraints, given by expression (9), while performing substitutions. Moreover, they must not re-assign more than v_{at}^q number of flights during time period t under scenario q .

In order to show the benefits of allowing intra-airline substitutions we perform an experiment using the same dataset described in Section 3.3, and assuming the ground delay cost of each flight to be uniformly distributed between 0.5 and 1.5. The flight-specific unit costs reflect differences in operating cost, payload, fare, class mix, downstream connectivity, and other factors that may cause an airline to attach higher or lower priority to certain flights. The cost ratio between airborne and ground holding -- λ -- is set to 25; i.e. there is no airborne holding in the solution of the Mukherjee-Hansen model. Substitutions are performed, for 100 different realizations of ground delay cost for each flight, to minimize airlines-specific total ground delay cost. Figure 5 below presents the box-and-whisker plot summarizing delay cost reduction, by airline, after substitutions.

Airlines that operate larger numbers of flights, such as AAL, benefit more from substitutions. This is intuitive, because an airline with a larger set of flights has more substitution opportunities. Figure 5 also shows the variation in cost reduction for the 100 different realizations of flight-specific ground delay unit cost. The upper and lower edge of each box represents the upper and lower quartiles respectively, while the diamond depicts the two standard deviation interval around the mean. There is considerable variability in cost reduction for certain airlines such as DAL and CHQ. In the former case, the explanation is that one DAL flight was assigned delay by the DRGH model. Depending on the cost reduction, this flight is substituted with another flight with lower unit delay cost in the substitution phase.

After a round of cancellations and intra-airline substitutions, airlines submit their revised slot assignment to the FAA. Due to cancellations, there may be some slots that remain unutilized. This can happen if an airline is unable to substitute one of its own flights to a slot that is vacated by a flight that the airline has decided to cancel. Furthermore, the updated information may become available on evolving conditions – i.e. scenario tree branching – in the time that is taken by the entire process of intra-airline substitution. The Mukherjee-Hansen model can then be re-run with updated information on both fronts -- capacity and demand. This will allow re-assigning some flights to the vacant slots created by cancellations. A slight modification to the objective function, by assigning airline-specific weights that are proportional to the number of flights cancelled by different airlines, may produce solutions that are analogous to the present-day Compression algorithm which gives priority to an airline that has cancelled a flight while performing slot re-assignments.

5. CONCLUSIONS

In this paper we describe two recently developed stochastic optimization models for the single airport ground holding problem, and present a methodology on how these models can be used to plan ground delay programs in face of uncertainty in airport capacity. The first model – the Ball et al. (2003) model -- is static stochastic, which means the ground delay decisions are made once and are not revised based on updated information on the airport conditions. The second model – the Mukherjee-Hansen (2003) model -- is dynamic stochastic one in which ground delays decisions are adaptive to the information on evolving airport capacity condition. In order to use these models in practice, it is necessary to allow airlines to participate in the planning process and perform substitutions that improve their internal objectives. In this paper we show how the solutions obtained from the stochastic models for the SAGHP can be used under the CDM paradigm.

The Ball et al. (2003) static stochastic model for the SAGHP produces the optimum set of PAARs that can be used by the FAA to set the airport acceptance rates when a GDP is implemented. Thereafter, an equitable slot allocation mechanism may be adopted to assign arrival slots to airlines, who can then perform substitutions and cancellations. After a round of intra-airline slot re-allocation, the FAA can run the Compression algorithm in order to utilize vacant slots that may have been created by cancellations. Thus, the Ball et al. model produces solutions that are easily adoptable in the current system. Their model, although formulated as an integer programming model, can be solved as an LP because the constraint matrix is totally unimodular. This eases the computation time, and thus making the model applicable in real time.

The Mukherjee-Hansen (2003) model, on the other hand, assigns scenario-contingent slots to individual flights. This however, does not restrict airlines to perform substitutions and cancellations. The solution to the Mukherjee-Hansen, although assigning slots to flights, may in practice be used as a means of allocating slots to airlines. Airlines can then perform substitutions among their flights while abiding by certain constraints discussed in this paper. Our experimental results show significant benefit, in terms of delay cost reduction, to airlines from this process.

Both models require as input the information on capacity scenarios and their probabilities. The Mukherjee-Hansen dynamic model also requires a probabilistic scenarios tree whose branching reveals unfolding conditions at the airport. Constructing capacity scenarios and scenario trees pose a challenge to the implementation of these models in practice. Recent work by Liu et al. (2006) shed some light in this direction. The research on stochastic models for the SAGHP, integrated with that on developing capacity scenarios and scenario tree, may lead to an efficient decision support tool in planning ground delay programs in face of uncertainty in the future.

6. REFERENCES

Andreatta, G., Brunetta, L., and Guastalla, G. (2000) 'From Ground Holding to Free Flight: An Exact Approach', *Transportation Science*, 34(4), pp. 394-401.

Ball, M.O., Hoffman, R., Odoni, A.R., and Rifkin, R. (2003) 'A Stochastic Integer Program with Dual Network Structure and Its Application to the Ground-Holding Problem', *Operations Research*, **51**(1), pp. 167-171.

Hoffman, R. (1997) *Integer Programming Models for Ground Holding in Air Traffic Flow Management*, Ph.D. Dissertation, University of Maryland, College Park, Maryland, U.S.A.

Hoffman, R. and Ball, M. O. (2000) 'A Comparison of Formulations for the Single-Airport Ground-Holding Problem with Banking Constraints', *Operations Research*, **48**(4), pp. 578-590.

Kotnyek, B., and Richetta, O. (2006) 'Equitable Models for the Stochastic Ground-Holding Problem Under Collaborative Decision Making', *Transportation Science*, **40**(2), pp. 133-146.

Liu, B., M. Hansen, A. Mukherjee. 2006. Scenario-Based Air Traffic Flow Management: Developing and Using Capacity Scenario Trees. *Transportation Research Record*, 1951, pp. 113-121.

Mukherjee, A., and Hansen, M. (2003) 'A Dynamic Stochastic Model for the Single Airport Ground Holding Problem', Working Paper, Berkeley, California, U.S.A.

Mukherjee, A. (2004) *Dynamic Stochastic Optimization Models for Air Traffic Flow Management*, Ph.D. Dissertation, University of California, Berkeley, California, U.S.A.

Odoni, A. R. (1987) 'The Flow Management Problem in Air Traffic Control', Odoni, A. R., Bianco, L., and Szego G. (ed) *Flow Control of Congested Networks*, Springer-Verlag, Berlin.

Richetta, O., and Odoni, A.R. (1993) 'Solving Optimally the Static Ground Holding Policy Problem in Air Traffic Control', *Transportation Science*, **27**(3), pp. 228-238.

Richetta, O., and Odoni, A.R. (1994) 'Dynamic Solution to the Ground Holding Problem in Air Traffic Control', *Transportation Research – Part A*, 28(3), pp. 167-185.

Terrab, M. and Odoni, A. R. (1993) 'Strategic Flow Management in Air Traffic Control', *Operations Research*, 41(1), pp. 138-152.

Vossen, T. (2002) *Fair Allocation Methods in Air Traffic Management*, Ph.D. Dissertation, University of Maryland, College Park, Maryland, U.S.A.

Table 1: Expected Delays (in aircraft-minutes) Obtained from the Stochastic Models

		Ground Delay	Airborne Delay	Total Delay Cost
$\lambda = 3$	Ball et al.	1170	336	2178
	Mukherjee-Hansen	1665	105	1980
$\lambda = 25$	Ball et al.	2835	0	2835
	Mukherjee-Hansen	2002	0	2002

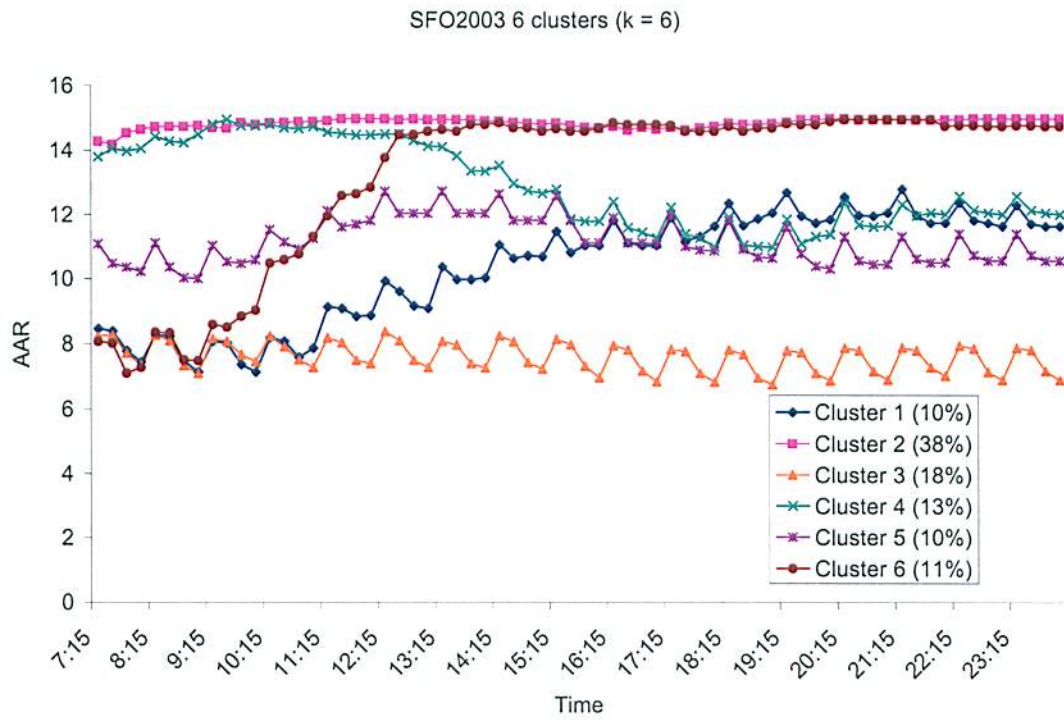


Figure 1: Capacity Scenarios and their Likelihood of Occurrence at SFO (source: Liu et al., 2006)

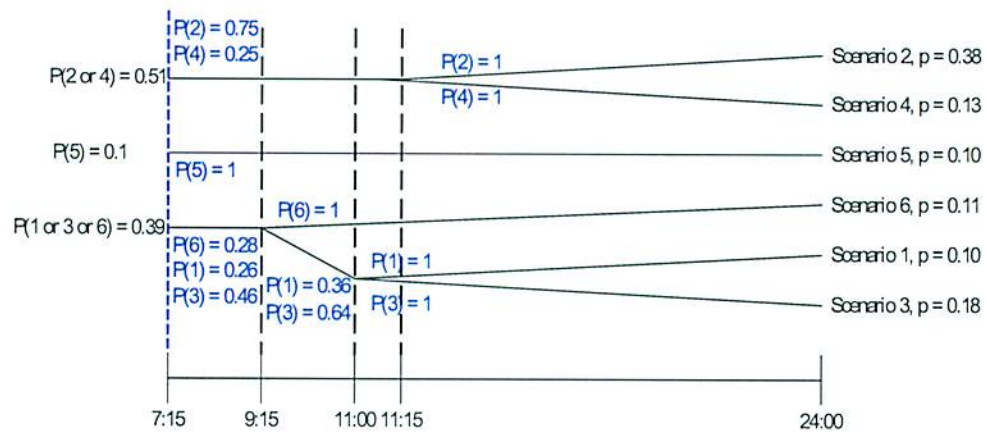


Figure 2: Capacity Scenario Tree and the Marginal Probabilities of the Scenarios

(source: Liu et al., 2006)

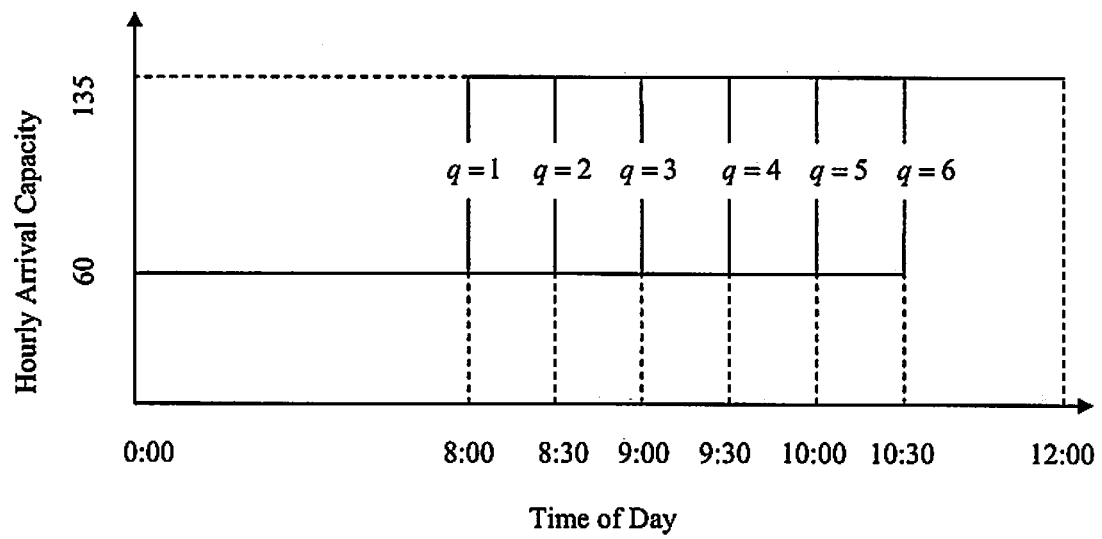


Figure 3: Capacity Scenarios for the Hypothetical Problem

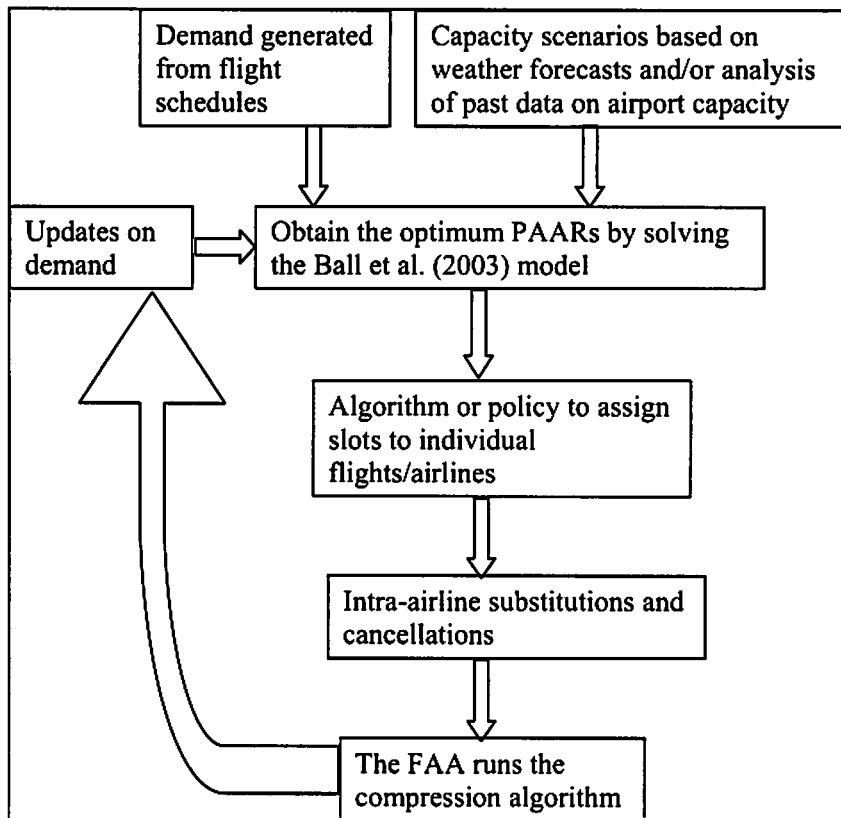


Figure 4: Collaborative GDP Planning Using the Ball et al. (2003) Model for the SAGHP

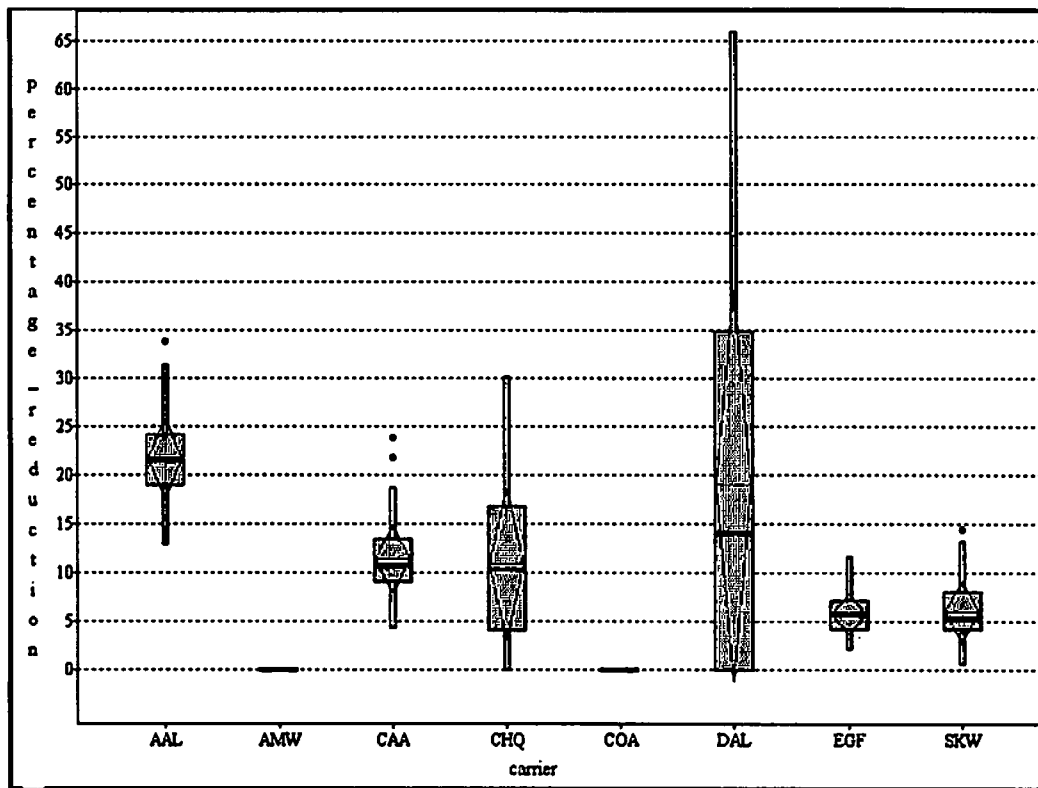


Figure 5: Box-and-Whisker Plot Showing Percentage Reductions in Ground Delay Costs of Different Airlines